

IN THE CLAIMS:

- 1 1. (Previously Presented) A system to move data from a source system to a destina-
2 tion system, comprising:
3 the source system in communication with the destination system, the source sys-
4 tem having a replication agent executing thereon, the replication agent adapted to gener-
5 ate a data stream, the data stream including:
6 a plurality of standalone headers, each of the plurality of standalone headers being
7 representative of a plurality of a data stream characteristics;
8 a data following header, the data following header including an extended attribute
9 field that associates an extended attribute with a data set information;
10 the destination system adapted to receive the data following header with the extended at-
11 tribute; and
12 an entry in a hidden permanent metadirectory of the destination system to store the data
13 set information associated with the extended attribute, the entry having the extended at-
14 tribute associated therewith so that retrieval of the entry from the hidden permanent
15 metadirectory also retrieves the extended attribute.
- 1 2. (Previously Presented) The system of claim 1, further comprising:
2 the plurality of standalone headers each include an indication of one of a plurality
3 of specialized header types and at least some of the plurality of specialized header types
4 are adapted for carrying directory inode data.
- 1 3. (Previously Presented) The system of claim 1, further comprising:
2 the data stream is adapted to carry source file system inode data and source file
3 generation numbers.

- 1 4. (Previously Presented) The system of claim 1, further comprising:
2 one of the standalone headers is a deleted files type and the directory inode data
3 includes a list of deleted files on the source file system.
- 1 5. (Previously Presented) The system of claim 1, further comprising:
2 the extended attributes include access control lists (ACLs) and streams associated
3 with a plurality of operating systems and system architectures.
- 1 6. (Previously Presented) The system of claim 1, further comprising:
2 one of the plurality of standalone headers includes an open file/undo header that
3 instructs the destination system to revert to an earlier copy of a stored file identified by
4 the open file/undo header.
- 1 7. (Previously Presented) The system of claim 1, further comprising:
2 the data set information includes file information.
- 1 8. (Previously Presented) The system of 1, further comprising:
2 the data set information includes changed files on the source system transmitted
3 for backup on the replica of the destination system.
- 1 9. (Previously Presented) The system of claim 1, further comprising:
2 the data following header includes offset and block number information with re-
3 spect to the data set information that follows the data following header.
- 1 10. (Previously Presented) The system of claim 1, further comprising:
2 the data following header includes a fixed-length record including (A) a generic
3 part for storing an indication of a data following header type; (B) a non-generic part,
4 adapted to carry predetermined data related to the extended attribute and data related to
5 offsets and block numbers for the data set information that follows the data following

6 header; and (C) a space for a bit-code representative of a name associated with the ex-
7 tended attribute.

1 11. (Previously Presented) The system of claim 1, further comprising:
2 each of the plurality of standalone headers includes a fixed-length record includ-
3 ing a generic part for storing an indication of one of a plurality of specialized header
4 types, a non-generic part, adapted to carry predetermined data related one of the special-
5 ized header types and a space for additional information.

1 12. (Previously Presented) The system of claim 1, further comprising:
2 the data following header is adapted to be positioned within the data stream at
3 predetermined intervals that are up to approximately 2 MB of data set information in size.

1 13. (Cancelled)

1 14. (Previously Presented) The system of claim 1, further comprising:
2 the destination system includes a hidden purgatory metadirectory in which current
3 data set information from the hidden permanent metadirectory is stored during an update
4 of the hidden permanent metadirectory with changed data set information, the destination
5 system being further adapted to (A) delete the hidden purgatory metadirectory after a
6 complete receipt of all expected changed data set information of the hidden permanent
7 metadirectory with the changed data set information, and (B) move current data set in-
8 formation stored on the hidden purgatory metadirectory back to the hidden permanent
9 metadirectory after an incomplete receipt of all expected changed data set information.

1 15. (Previously Presented) The system of claim 14, further comprising:
2 the destination system is adapted to create a hidden new metadirectory to store
3 changed data set information for transfer to the hidden permanent metadirectory after of
4 the complete receipt of all the expected changed data set information.

1 16. (Previously Presented) The system of claim 1, further comprising:
2 the source system and the destination system are remote with respect to each other
3 and interconnected by a network, and wherein the data stream is encapsulated within a
4 networking protocol adapted for transmission over the network.

1 17. (Previously Presented) A system for transmitting a data stream that includes data
2 set information, the system comprising:
3 a destination system having a replica stored thereon;
4 a source system having a replication agent executing thereon, the replication agent
5 adapted to generate the data stream, the data stream having a data following header ap-
6 pended to a predetermined-sized chunk of data, the data following header including a
7 field that identifies extended attributes associated with data set information carried in the
8 chunk; and
9 the destination system, in response to the data following header, storing the chunk of data
10 as an entry in a hidden permanent metadirectory in a file system of the destination sys-
11 tem, the entry having the extended attribute associated therewith so that retrieval of the
12 entry from the hidden permanent metadirectory also retrieves the extended attribute.

1 18. (Previously Presented) The system of claim 17, further comprising:
2 the extended attributes include access control lists (ACLs) and streams associated
3 with a plurality of operating systems and system architectures.

1 19. (Previously Presented) The system of claim 17, further comprising:
2 the data set information includes file information.

1 20. (Previously Presented) The system of claim 17, further comprising:
2 the data set information includes changed files on the source system transmitted
3 for backup on the replica of the destination system.

- 1 21. (Previously Presented) The system of claim 17, further comprising:
2 the data following header includes offset and block number information with re-
3 spect to the data set information that follows the data following header.
- 1 22. (Previously Presented) The system of claim 17, further comprising:
2 the data following header includes a fixed-length record including (A) a generic
3 part for storing an indication of a data following header type; (B) a non-generic part,
4 adapted to carry predetermined data related to the extended attribute and data related to
5 offsets and block numbers for the data set information that follows the data following
6 header; and (C) a space for a bit-code representative of a name associated with the ex-
7 tended attribute.
- 1 23. (Previously Presented) The system of ~~in~~ claim 17, further comprising:
2 the chunk has a size of up to approximately 2 MB of data set information.
- 1 24. (Previously Presented) The system of claim 17, further comprising:
2 the destination system is adapted to receive the data following header with the
3 extended attribute and cause the data set information associated with the extended attrib-
4 ute to be stored in an entry in a hidden permanent metadirectory with identifiers that are
5 the same as identifiers for the data set information in a file system of the destination sys-
6 tem, the entry having the extended attribute associated therewith so that retrieval of the
7 entry from the hidden permanent metadirectory also retrieves the extended attribute.
- 1 25. (Previously Presented) The system of claim 24, further comprising:
2 the destination system includes a hidden purgatory metadirectory in which current
3 data set information from the hidden permanent directory is stored during an update of
4 the hidden permanent metadirectory with changed data set information, the destination
5 system being further adapted to (A) delete the hidden purgatory metadirectory after a

6 complete receipt of all expected changed data set information of the hidden permanent
7 metadirectory with the changed data set information, and (B) move current data set in-
8 formation stored on the hidden purgatory metadirectory back to the hidden permanent
9 metadirectory after an incomplete receipt of all expected changed data set information.

1 26. (Previously Presented) The system of claim 25, further comprising:
2 the destination system is adapted to create a hidden new metadirectory to store
3 changed data set information for transfer to the hidden permanent metadirectory after of
4 the complete receipt of all the expected changed data set information.

1 27. (Previously Presented) A method for storing and retrieving extended attributes
2 associated with a data set information comprising:
3 executing a replication agent on a source system, the replication agent adapted to gener-
4 ate a data stream, the data stream having;
5 a plurality of standalone headers, each of the plurality of standalone headers being
6 representative of a plurality of data stream characteristics;
7 a data following header, the data following header including an extended attribute
8 field that associates an extended attribute with a chunk of data preceding the data
9 following header;
10 receiving the data stream by a destination system;
11 storing in a destination system, in response to receiving the data stream, a current
12 data set information with current extended attributes in a permanent hidden metadirec-
13 tory;
14 transferring the data set information to a purgatory metadirectory upon receipt of
15 a changed data set information;
16 storing the received changed data set information in a new metadirectory; and
17 upon completion of receipt of all expected changed data set information, transfer-
18 ring the received changed data set information from the new metadirectory to the perma-

19 nent metadirectory, the permanent metadirectory thereby being available for retrieval of
20 extended attributes associated with the data set information.

1 28. (Previously Presented) The method of claim 27, further comprising:
2 upon a failure to complete receipt of all expected changed data set information,
3 transferring the current data set information with the current extended attributes back to
4 the permanent metadirectory.

1 29. (Previously Presented) The method of claim 28, further comprising:
2 the data set information including files organized in a directory tree structure the
3 same as a file system structure on the destination system and the extended attributes in-
4 cluding access control lists (ACLs) and streams associated with the files.

1 30. (Previously Presented) The method of claim 29, further comprising:
2 deleting the purgatory metadirectory after one of, either (A) the transferring of the
3 changed data set information from the new metadirectory to the permanent metadirectory
4 or (B) the transferring of the current data set information from the purgatory metadirec-
5 tory back to the permanent metadirectory.

1 31. (Previously Presented) The method of claim 30, further comprising:
2 upon a request from the source to restore data sets from the data set information,
3 scanning the permanent directory and retrieving the data sets including retrieving respec-
4 tive of the extended attributes associated with the data sets.

1 32. (Previously Presented) The method of claim 31, further comprising:
2 providing the retrieved data sets' extended attributes in a format for transmission
3 to the source from the destination, the format including data following headers each hav-
4 ing a field that associates the respective of the extended attributes with the retrieved data
5 sets.

- 1 33. (Previously Presented) The method of claim 32, further comprising:
2 associating the respective of the extended attributes with the data sets based upon
3 NT streams.
- 1 34. (Previously Presented) The method of claim 27, further comprising:
2 associating the extended attributes with the data set information in the permanent
3 metadirectory using NT streams.
- 1 35. (Previously Presented) The method of claim 27, further comprising:
2 providing the data sets' extended attributes in a format for transmission to the des-
3 tination from the source, the format including data following headers each having a field
4 that associates the respective of the extended attributes with the retrieved data sets.
- 1 36. (Cancelled)
- 1 37. (Cancelled)
- 1 38. (Previously Presented) A computer readable medium for storing and retrieving
2 extended attributes associated with a data set information, the computer readable medium
3 including program instructions for performing the steps of:
4 executing a replication agent on a source system, the replication agent adapted to gener-
5 ate a data stream, the data stream having:
6 a plurality of standalone headers, each of the plurality of standalone headers being
7 representative of a plurality of data stream characteristics;
8 a data following header, the data following header including an extended attribute
9 field that associates an extended attribute with a chunk of data preceding the data
10 following header;

11 receiving the data stream by a destination system;
12 storing in the destination system, in response to receiving the data stream, a cur-
13 rent data set information with current extended attributes in a permanent hidden metadi-
14 rectory;
15 transferring the data set information to a purgatory metadirectory upon receipt of
16 a changed data set information;
17 storing the received changed data set information in a new metadirectory; and
18 upon completion of receipt of all expected changed data set information, transfer-
19 ring the received changed data set information from the new metadirectory to the perma-
20 nent metadirectory, the permanent metadirectory thereby being available for retrieval of
21 extended attributes associated with the data set information.

1 39. (Cancelled)

1 40. (Previously Presented) A system for transmitting a data stream, the system com-
2 prising:
3 an operating system for a destination data storage system, the operating system
4 configured to receive the data stream and maintain a replica of a data set of the source
5 system; having a replica residing therein;
6 a plurality of standalone headers in the data stream, the plurality of standalone
7 headers having discrete identifiers, each of the plurality of standalone headers being rep-
8 resentative of a plurality of data stream characteristics;
9 a data following header in the data stream, the data following header having an
10 extended attribute field that associates an extended attribute with the data in the data
11 stream, the extended attribute including metadata external to the data set; and
12 the destination system, in response to the data following header, storing the data set in-
13 formation associated with the extended attribute as an entry in a hidden permanent meta-
14 directory with identifiers that are the same as identifiers for the data set information in a
15 file system of the destination system, the entry having the extended attribute associated

16 therewith so that retrieval of the entry from the hidden permanent metadirectory also re-
17 trieves the extended attribute.

1 41. (Previously Presented) A method for backing up data from a source to a destina-
2 tion data storage system, the method comprising:
3 formatting, by the source, a data stream including system specific attributes; by append-
4 ing; to the data stream a plurality of standalone headers and a plurality of data following
5 headers:
6 the standalone headers including information which characterize data in
7 the data stream;
8 the data following headers including at least one extended attribute which identi-
9 fies a system specific attribute of data in the data stream; ;
10 communicating the formatted data stream from the source to the destination data storage
11 system; and
12 storing, by the destination system, the data with the system specific attribute as an entry
13 in a hidden permanent metadirectory with identifiers that are the same as identifiers for
14 the data in a file system of the destination system, the entry having the system specific
15 attribute associated therewith so that retrieval of the entry from the hidden permanent
16 metadirectory also retrieves the system specific attribute.

1 42. (Previously Presented) The method of claim 41, further comprising:
2 interpreting the information, by the destination data storage system, as instructions for
3 handling data in the data stream according to the system specific attribute if the system
4 specific attribute is compatible with the destination data storage system; and
5 interpreting the information, by the destination data storage system, as instructions for
6 ignoring the system specific attribute if the system specific attribute is incompatible with
7 the destination data storage system

- 1 43. (Previously Presented) The method of claim 41, further comprising:
2 including space in the data following header for the extended attributes to store
3 self-descriptive information.
- 1 44. (Previously Presented) A method for storing data in a destination system, the data
2 obtained from a source system, comprising:
3 executing a replication agent on the source system, the replication agent adapted
4 to generate a data stream, the data stream having, a plurality of standalone headers and a
5 data following header, the data following header including an extended attribute field that
6 associates an extended attribute with a chunk of data associated with the data following
7 header;
8 receiving the data stream by the destination system;
9 storing in the destination system, in response to receiving the data stream, the
10 chunk of data associated with the data following header as an entry in a hidden metadi-
11 rectory; and
12 retrieving the chunk of data from the hidden metadirectory so that retrieval of the
13 entry from the hidden metadirectory also retrieves the extended attribute.
- 1 45. (Previously Presented) The method of claim 44, further comprising:
2 transferring the chunk of data to a purgatory metadirectory upon receipt of a
3 changed chunk of data;
4 storing the received changed chunk of data in a new metadirectory;
5 storing additional changed chunks of data in the new metadirectory; and
6 upon completion of receipt of all expected additional changed chunks of data,
7 transferring the received changed chunks of data from the new metadirectory to a perma-
8 nent metadirectory, the permanent metadirectory then being available for retrieval of ex-
9 tended attributes associated with the received changed chunks of data.

1 46. (Previously Presented) The method of claim 44, further comprising:
2 including in the standalone headers a deleted files type, and the data including a
3 list of deleted files on the source system.

1 47. (Previously Presented) The method of claim 44, further comprising:
2 including access control lists (ACLs), streams associated with a plurality of oper-
3 ating systems, and system architectures in the extended attributes.

1 48. (Previously Presented) The method of claim 44, further comprising:
2 including as one of the standalone headers an open file/undo header that instructs
3 the destination system to revert to an earlier copy of a stored file identified by the open
4 file/undo header.

1 49. (Previously Presented) The method of claim 44, further comprising:
2 including changed files on the source system, the changed files transmitted for
3 backup on the replica.

1 50. (Previously Presented) The method of claim 44, further comprising:
2 including offset and block number information in a header for a data set that
3 follows the header.

1 51. (Previously Presented) The method of claim 44, further comprising:
2 including in the data following header a record having (A) a generic part for stor-
3 ing an indication of a data following header type; (B) a non-generic part, adapted to carry
4 predetermined data related to the extended attribute and data related to offsets and block
5 numbers for a data set information that follows the data following header; and (C) a space
6 for a bit-code representative of a name associated with the extended attribute.

1 52. (Previously Presented) The method of claim 44, further comprising:
2 including in at least one of the plurality of standalone headers a fixed-length re-
3 cord including a generic part for storing an indication of one of a plurality of specialized

4 header types, a non-generic part, adapted to carry predetermined data related one of the
5 specialized header types and a space for additional information.

1 53. (Previously Presented) The method of claim 44, further comprising:
2 adapting the data following header to be positioned within the data stream at pre-
3 determined intervals that are up to approximately 2 MB of data set information in size.

1 54. (Previously Presented) The method of claim 44, further comprising:
2 including in the destination system a hidden purgatory metadirectory in which
3 current data set information from the hidden permanent metadirectory is stored during an
4 update of the hidden permanent metadirectory with changed data set information, the des-
5 tination system being further adapted to (A) delete the hidden purgatory metadirectory
6 after a complete receipt of all expected changed data set information of the hidden per-
7 manent metadirectory with the changed data set information, and (B) move current data
8 set information stored on the hidden purgatory metadirectory back to the hidden perma-
9 nent metadirectory after an incomplete receipt of all expected changed data set informa-
10 tion.

1 55. (Previously Presented) The method of claim 44, further comprising:
2 adapting the destination system to create a hidden new metadirectory to store
3 changed data set information for transfer to the hidden permanent metadirectory after re-
4 ceipt of all the expected changed data set information.

1 56. (Previously Presented) The method of claim 44, further comprising:
2 interconnecting the source system and the destination system by a computer net-
3 work, and the data stream is encapsulated within a networking protocol adapted for
4 transmission over the network.

1 57. (Previously Presented) The method of claim 44, further comprising:
2 carrying source file system inode data and source file generation numbers in the
3 data stream.

1 58. (Previously Presented) A system to store data in a destination system, the data ob-
2 tained from a source system, comprising:
3 a replication agent to execute on the source system, the replication agent adapted
4 to generate a data stream, the data stream having, a plurality of standalone headers and a
5 data following header, the data following header including an extended attribute field that
6 associates an extended attribute with a chunk of data associated with the data following
7 header;
8 a port to receive the data stream by the destination system;
9 an entry in a hidden metadirectory in the destination system to store, in response
10 to receiving the data stream, the chunk of data associated with the data following header;
11 and
12 retrieving the chunk of data from the hidden metadirectory so that retrieval of the
13 entry from the hidden metadirectory also retrieves the extended attribute.

1 59. (Previously Presented) The system of claim 58, further comprising:
2 transferring the chunk of data to a purgatory metadirectory upon receipt of a
3 changed chunk of data;
4 storing the received changed chunk of data in a new metadirectory;
5 storing additional changed chunks of data in the new metadirectory; and
6 upon completion of receipt of all expected additional changed chunks of data,
7 transferring the received changed chunks of data from the new metadirectory to a perma-
8 nent metadirectory, the permanent metadirectory then being available for retrieval of ex-
9 tended attributes associated with the received changed chunks of data.

1 60. (Previously Presented) The system of claim 58, further comprising:
2 a deleted files type included in the standalone header, and the data including a list
3 of deleted files on the source system.

1 61. (Previously Presented) The system of claim 58, further comprising:
2 access control lists (ACLs), streams associated with a plurality of operating sys-
3 tems, and system architectures included in the extended attributes.

- 1 62. (Previously Presented) The system of claim 58, further comprising:
2 an open file/undo header that instructs the destination system to revert to an ear-
3 lier copy of a stored file identified by the open file/undo header included as one of the
4 standalone headers.
- 1 63. (Previously Presented) The system of claim 58, further comprising:
2 changed files on the source system, the changed files transmitted for backup on
3 the replica.
- 1 64. (Previously Presented) The system of claim 58, further comprising:
2 offset and block number information in included in a header for a data set that fol-
3 lows the header.
- 1 65. (Previously Presented) The system of claim 58, further comprising:
2 a record having (A) a generic part for storing an indication of a data following
3 header type; (B) a non-generic part, adapted to carry predetermined data related to the
4 extended attribute and data related to offsets and block numbers for a data set information
5 that follows the data following header; and (C) a space for a bit-code representative of a
6 name associated with the extended attribute included in the data following header.
- 1 66. (Previously Presented) The system of claim 58, further comprising:
2 a fixed-length record including a generic part for storing an indication of one of a
3 plurality of specialized header types, a non-generic part, adapted to carry predetermined
4 data related one of the specialized header types, and a space for additional information
5 included in at least one of the plurality of standalone headers.
- 1 67. (Previously Presented) The system of claim 58, further comprising:
2 a processor to position the data following header to be positioned within the data
3 stream at predetermined intervals that are up to approximately 2 MB of data set informa-
4 tion in size.
- 1 68. (Previously Presented) The system of claim 58, further comprising:

2 a hidden purgatory metadirectory included in the destination system, the hidden
3 purgatory metadirectory storing current data set information from the hidden permanent
4 metadirectory during an update of the hidden permanent metadirectory with changed data
5 set information, the destination system being further adapted to (A) delete the hidden pur-
6 gatory metadirectory after a complete receipt of all expected changed data set information
7 of the hidden permanent metadirectory with the changed data set information, and (B)
8 move current data set information stored on the hidden purgatory metadirectory back to
9 the hidden permanent metadirectory after an incomplete receipt of all expected changed
10 data set information.

1 69. (Previously Presented) The system of claim 58, further comprising:

2 a hidden new metadirectory created in the destination system to store changed
3 data set information for transfer to the hidden permanent metadirectory after receipt of all
4 the expected changed data set information.

1 70. (Previously Presented) The system of claim 58, further comprising:

2 a computer network to interconnect the source system and the destination system,
3 and the data stream is encapsulated within a networking protocol adapted for transmis-
4 sion over the network.

1 71. (Previously Presented) The system of claim 58, further comprising:

2 source file system inode data and source file generation numbers carried in the
3 data stream.

1 72. (Previously Presented) A computer readable media, comprising:

2 said computer readable media containing instructions for execution on a processor
3 for a method of storing data in a destination system, the data obtained from a source sys-
4 tem, the method having the steps of:

5 executing a replication agent on the source system, the replication agent
6 adapted to generate a data stream, the data stream having, a plurality of stand-
7 alone headers and a data following header, the data following header including an

8 extended attribute field that associates an extended attribute with a chunk of data
9 associated with the data following header;
10 receiving the data stream by the destination system;
11 storing in the destination system, in response to receiving the data stream,
12 the chunk of data associated with the data following header as an entry in a hid-
13 den metadirectory; and
14 retrieving the chunk of data from the hidden metadirectory so that retrieval
15 of the entry from the hidden metadirectory also retrieves the extended attribute.